

## Facilitating High-Throughput ASC Calculations


John T. Daly, HPC-4

The productivity of the application user is arguably one of the most significant metrics upon which high-performance computing (HPC) systems performance is evaluated for the Advanced Simulation and Computing (ASC) program. One important component of that productivity is the application throughput associated with the calculation. Application throughput for *capacity* calculations, characterized by large numbers of jobs on relatively few processors, is typically limited by the performance and scaling of an HPC system. However, in the case of *capability* calculations, characterized by relatively few jobs run on a significant fraction of the platform, many thousands of processors, the application throughput, or movement of data, is frequently limited by the reliability and availability of the system [1,2].

This paper discusses improvements in the throughput of HPC applications on capability systems realized by developing and implementing a unique *job\_suite* toolkit.

The *job\_suite* toolkit—which includes *job\_submit*, *job\_view*, *job\_list*, and *job\_kill*—was developed by the author specifically for the purpose of keeping these capability class calculations running and maximizing system utilization in the face of various types of interrupts. This is accomplished by a combination of job monitoring, to check that a running application is making progress, and automated restart, to get a stopped application running again after an interrupt. The *job\_suite* tools have run successfully on ASC platforms including Red Storm, Purple, BG/L, and Q. They can interact with jobs via any of the tri-Lab-supported schedulers

**Fig. 1.**  
Progress on the W76  
LEP calculations in  
terms of operational  
utilization over time.



(i.e., LSF, LCRM, PBS, or Moab) or through a Posix interface on platforms without a scheduler.


The *job\_suite* set of tools has been used to facilitate user productivity for other calculations [3], but the most compelling example of its effectiveness comes from the W76 LEP calculation run in tandem by a single user on Red Storm, Purple, and BG/L between December 2006 and February 2007. This calculation consisted of seven high-priority jobs, each utilizing 4096 to 8192 processors. It was consistently run on one-quarter or more of the available compute resources of all three platforms. The jobs typically sustained an aggregate of 11 interrupts per day. Figure 1 illustrates the operational utilization over time of the seven different jobs, and Fig. 2 depicts the wall-clock time accumulated by the jobs in processor hours. This level of throughput in face of frequent interrupts represents a substantial HPC capability for the weapons program.

- [1] J.T. Daly, "Methodology and Metrics for Quantifying Application Throughput," Proceedings of the Nuclear Explosives Code Developers' Conference (NECDC), in press.
- [2] J.T. Daly, "A Higher Order Estimate of the Optimum Checkpoint Interval for Restart Dumps," *Future Generation Computer Systems* **22**, 300 (2006).
- [3] J.T. Daly, "CIS External Review: A Customer Perspective on Red Storm," Los Alamos Technical Report, LA-UR-06-3802.

*For more information contact John Daly at [jtd@lanl.gov](mailto:jtd@lanl.gov).*

#### Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.



**Fig. 2.**  
Progress on the W76 LEP calculations in terms of accumulated wall-clock hours of CPU-core time.